EXPRESS MAIL NO. 73490258US

PATENT
DOCKET NO. 00-8005

UNITED STATES PATENT APPLICATION

OF

Siva Perraju TOLETY

FOR

METHOD, APPARATUS AND PROGRAM FOR DETERMINING AVAILABLE BANDWIDTH BETWEEN MULTIPLE POINTS IN A COMMUNICATION SYSTEM

BACKGROUND OF THE INVENTION

1. Field of the Invention

This invention relates generally to communication systems, and in particular to a method, apparatus and program for determining available bandwidth in a communication path coupled between nodes in a communication system.

2. Related Art

Internet connection service providers typically promise customers that they will be provided with specific bandwidth rates for particular types of service connections (e.g., ADSL, xDSL, ISDN, etc.). Despite these promises, however, at any given time, actual bandwidth rates may differ substantially from promised bandwidth rates, owing to, for example, the presence of severe traffic congestion in communication system components and system component capacity limitations. As a result, connection service providers often receive many complaints from customers concerning long download delays, problems encountered during Packet Internet Groper (PING) operations, and other complaints relating to bandwidth reductions in general.

To respond to these problems, connection service providers often employ known test procedures for isolating problem system components. Unfortunately,

most known test procedures are unsatisfactory in that they test only components interposed between customer premise and central office switching equipment, but do not test upstream system components. An example of one such test procedure is the Fujitsu Speed Port Shelf Manager, which enables troubleshooters to conduct bit error rate (BER) tests, noise margin estimates, errored seconds (ER) and severe errored seconds (SER) estimates, and attenuation estimates.

At least some connection service providers respond to customer complaints concerning low bandwidth rates by terminating existing virtual circuits connecting customer premise equipment to backbone cloud (network) equipment, such as a network (e.g., Frame Relay) switch, and by then "rebuilding" other virtual circuits to couple the customer premise equipment to a test server through another switch in the network. This step is necessary because, during normal, non-testing conditions, the test server and customer premise equipment are typically connected to different switches, and thus are not communicatively coupled together. A file having a predetermined size is then downloaded from the test server to the customer premise equipment by way of the rebuilt virtual circuit. The customer premise equipment then measures (using a program) the period of time taken for the

file to be received therein, by, for example, detecting receipt times of beginning and ending portions (e.g., Start-of-File and End-of-File, respectively) of the file, and by then calculating the difference between those receipt times. The customer premise equipment also determines the size of the downloaded file by counting each byte included in the file, as it is received in the customer premise equipment, and by then multiplying the total number of counted bytes by '8' to determine the total number of bits included in the file. Thereafter, an estimate is made of the file download rate (i.e., the downlink bandwidth rate), based on the measured download time period and the determined file size.

Unfortunately, however, the foregoing prior art test procedure has a number of drawbacks. One drawback is that the " rebuilt" virtual circuit is not necessarily the same original virtual circuit used during normal, non-testing conditions, and thus the bandwidth rate determined during the test may be an inaccurate estimation of the typical bandwidth provided to the customer premise equipment. Also, the test procedure does not provide any indication of the uplink bandwidth rate and the system components which may be causing the bandwidth reduction problem. Moreover, the test procedure requires intensive operator

intervention for rebuilding the virtual circuit, rendering the procedure susceptible to human-induced errors. Furthermore, owing to possible manpower limitations and associated costs, it might not be feasible to perform such a procedure on a large scale, especially where the customer base being supported is a large one.

At least one known bandwidth estimation technique enables customers to conduct a download bandwidth test using off-the-shelf software (see, for example, the " Bandwidth Speed Test" provided at http://www.computingcentral.com/topics/bandwidth/speedtest500.asp). This technique apparently is performed using a large Hyper Text Markup Language (HTML) page, wherein a Javascript code in the page determines starting and ending times of a transfer of a portion of the page, for use in determining the download bandwidth. Unfortunately, because browser software is employed to determine the starting and ending times of the page transfer, the technique is subject to browser idiosyncrasies which can reduce the accuracy of the downlink bandwidth determination (different browser software may provide different bandwidth estimates). The technique also requires the use of customer premise equipment software which can understand the HyperText Transfer Protocol (HTTP). This can further

reduce the accuracy of the bandwidth determination, especially in cases where the customer premise equipment software gives low priority to processing HTTP-related requests. Moreover, the technique does not provide any estimate of the uplink bandwidth.

There is a need, therefore, for an improved technique which reliably determines an amount of bandwidth available in a communication path coupling together nodes in a communication system, and does not suffer from the drawbacks discussed above.

SUMMARY OF INVENTION

It is a first object of this invention to provide an improved method, apparatus, and program for determining an amount of bandwidth available in at least a portion of a communication path coupling together nodes in a communication system.

It is another object of this invention to determine an amount of bandwidth available between multiple points in a communication system, at a single node in the communication system.

Further objects and advantages of this invention will become apparent from a consideration of the drawings and ensuing description.

The foregoing and other problems are overcome and the objects of the invention are realized by a method for determining an amount of bandwidth available in at least one communication path which couples a plurality of nodes together, and a program and apparatus that operate in accordance with that method. In accordance with one embodiment of the invention, the method comprises steps of exercising the communication path, using information signals, to determine the amount of time it takes for at least one of those information signals to traverse the path in at least one direction, and determining the amount of bandwidth available in at least a portion of the path, based on the amount of time determined in the exercising step.

A first one of the plurality of nodes preferably comprises a router located at a Point of Presence of an Internet Service Provider (ISP), and a second one of the plurality of nodes preferably comprises a user communication terminal (customer premise equipment). Those nodes are coupled together through components of a communication system forming the communication path.

In accordance with another embodiment of this invention, uplink and downlink bandwidth rates available in the communication path are determined by transferring a file between a test node and the user

communication terminal, by way of the at least one communication path and a router. The uplink and downlink bandwidth rates are then calculated based on the file size, a rate at which the file is received at the terminal, and a rate at which the file is received at the test node, respectively.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be more readily understood from a detailed description of the preferred embodiments taken in conjunction with following figures:

Fig. 1 is a block diagram of a communication system 10 that includes a test node 22 constructed and operated according to this invention.

Fig. 2 is a block diagram of a user communication terminal 21 representing in further detail the test node 22 and a user communication terminal 1 of the system 10 of Fig. 1.

Figs. 3a-3b show a logical flow diagram of a method for determining an amount of bandwidth available in a communication path coupling multiple points in a communication system, in accordance with one embodiment of this invention.

Fig. 4 shows an example of a message format employed for communicating messages during the performance of the method of Figs. 3a-3b, according to one embodiment of the invention.

Fig. 5 shows an example of another message format employed for communicating error messages during the performance of the method of Figs. 3a-3b, according to an embodiment of the invention.

Fig. 6 is a block diagram of another communication system 50 that is suitable for practicing this invention, wherein the communication system 50 includes a test node 22 constructed and operated according to this invention.

Fig. 7 shows a logical flow diagram of a method for determining an amount of bandwidth available in a communication path coupled between multiple points in a communication system, according to another embodiment of this invention.

Identical portions of the various figures have been identified with the same reference numerals in order to simplify the description of the invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

Fig. 1 is a block diagram of a communication system 10 that is suitable for practicing this invention. In the illustrated embodiment, the communication system 10 comprises client premise equipment (CPE) 18, a central office switching station 8, a communication network 13, and a communication interface 20 for coupling the communication network 13 to a communication network entity 17, such as the Internet. In accordance with this invention, the communication system 10 also comprises a test node (hereinafter also referred to as a " user communication terminal") 22 which is bidirectionally coupled to a node 15 (to be described below) of the interface 20, through each of a plurality of communication links 24a and 24b.

The CPE 18 is bidirectionally coupled to transceiving equipment 7 of the central office switching station 8 through a communication interface 6, such as a telephone line (e.g., landline trunk), although in other embodiments, other suitable types of interfaces may also be employed for that interface 6, such as one or more coaxial cable lines, or a wireless interface. A multiplexer/demultiplexer device 9 of the central office switching station 8 is bidirectionally coupled to the network 13 through a communication interface

12, which, in the preferred embodiment, includes a T1 or T3 high speed link, although in other embodiments, other suitable types interfaces also may be employed between those components 9 and 13, such as, for example, a wireless or other interface, depending on applicable system architecture.

The test node 22 according to this invention includes a plurality of interfaces (IF1) and (IF2) that are each coupled to the node 15 of communication interface 20 through a respective one of the bidirectional communication links 24a and 24b. Preferably, the test node 22 includes a PC or a server computer, each interface (IF1) and (IF2) includes a network interface card (NIC1 and NIC2, respectively) having a unique, pre-assigned IP address (i.e., the test node 22 is a "multi-homed" device), and the links 24a and 24b each include a high-speed link, such as a T1 or T3 link. Preferably, the links 24a and 24b do not support any other traffic other than that provided between the node 22 and router 15, and thus each link 24a and 24b has a same, known available bandwidth capacity. The internal construction of the test node 22 and the manner in which that test node 22 is employed in the invention will be described in more detail below.

The CPE 18 shown in Fig. 1 will now be described. In the illustrated embodiment, the CPE 18 includes transceiving equipment 3 and one or more user communication terminals, such as, for example, a PC 1 and a telephone 2, that are bidirectionally coupled to the transceiving equipment 3 through respective communication interfaces 5 and 19. Preferably, the transceiving equipment 3 includes an Asynchronous Digital Subscriber Line (ADSL) modem, although in other embodiments, other suitable types of transceiving equipment may also be employed, such as, for example, an xDSL modem, an Integrated Services Digital Network (ISDN) modem, a cable modem (and/or a cable converter or set top box), wireless transceiving equipment, and the like, depending on, for example, applicable performance criteria and the types of user communication terminals 1 and 2 employed. In the preferred embodiment, the ADSL modem 3 operates by modulating voice signals and data received from the respective devices 2 and 1, using a known ADSL modulation technique (such as, e.g., Discrete Multitone Technology (DMT), Carrierless Amplitude Modulation (CAP), or Multiple Virtual Line (MVL) technique, etc.), and by forwarding resulting modulated information to the switching station 8 by way of interface 6. The modem 3 also operates by

demodulating information received over the interface 6, using a known ADSL demodulation technique (e.g., DMT, CAP, or MVL techniques, etc.), and by separating voice signals from data included in the received information using, for example, an associated splitting device (e.g., a POTS splitter) (not shown). The separated voice signals and data are then forwarded to the respective devices 2 and 1, through the corresponding interfaces 19 and 5, each of which may include, for example, a telephone line, cable line, or a wireless interface, depending on the types of user communication terminals 1 and 2 employed.

Referring now to Fig. 2, a preferred embodiment of the test node 22 (also referred to as a " user communication terminal"), and an exemplary embodiment of the user communication terminal 1, are shown, and are each identified by reference numeral 21. In Fig. 2, user communication terminal 21 preferably comprises a controller (e.g., a microprocessor and/or logic array) 21a for performing arithmetic and/or logical operations required for program execution, at least one input user-interface 21d that is coupled to the controller 21a, and at least one output user-interface 21e that also is coupled to the controller 21a. In the case of the user communication terminal 22 (i.e., the test node 22 of Fig. 1), the controller 21a is

bidirectionally coupled to each of the interfaces (IF1) and (IF2) (i.e., NIC1 and NIC2, respectively), and can communicate bidirectionally through each one of those interfaces and the corresponding link 24a, 24b (Fig. 1) coupled thereto. In the case of the user communication terminal 1 (of Fig. 1), one or more ports 21b are provided for enabling the controller 21a of that terminal 1 to communicate bidirectionally with the modem 3 through those port(s) 21b and the link 5.

The input user-interface 21d may include any suitable type of user-operable input device(s), such as, for example, a keyboard, mouse, touch screen, or trackball, and the output user-interface 21e may include, for example, a video display, a liquid crystal or other flat panel display, a printer, a speaker, and/or any other suitable type of output device for enabling a user to perceive outputted information. For the purposes of this description, the output user-interface 21e is assumed to be a display.

The user communication terminal 21 of Fig. 2 also includes at least one memory (e.g., disk drives, read-only memories, and/or random access memories) 21c that is bidirectionally coupled to the controller 21a. The memory 21c stores temporary data and instructions, and

also stores various routines and operating programs (e.g., Microsoft Windows, UNIX/LINUX, or OS/2) that are used by the controller 21a for controlling the overall operation of the user communication terminal 21. Preferably, at least one of the programs (e.g., Microsoft Winsock) stored in memory 21c adheres to TCP/IP protocols (i.e., includes a TCP/IP stack), for implementing a known method for connecting the terminal 21 to the Internet 17, through the various intermediate components of the system 10. The memory 21c may also store web browser software, such as, for example, Microsoft Internet Explorer (IE) and/or Netscape Navigator, for enabling a user of the terminal 21 to navigate or otherwise exchange information with the World Wide Web (WWW). The memory 21c of test node 22 also stores routines for implementing a method according to one embodiment of this invention. That method will be described below in relation to Figs. 3a-3c. In accordance with another embodiment this invention, both the user communication terminal 1 and the test node 22 store routines for implementing another method of the invention, which will be described below in relation to Fig. 7.

Before describing the further components of the communication system 1, it should be noted that

although this invention is described in the context of the test node 22 and user communication terminal 1 of Fig. 1 each being embodied as a PC (or, in the case of test node 22, a server computer), any other suitable type of user communication terminal may be employed for those devices 1 and 22, and the CPE 18 may include other user communication terminals, in addition to those depicted in Fig. 1. For example, in other embodiments, the individual devices 1 and 22 each may be embodied as a portable PC docking node, a web TV, personal digital assistant, handheld personal digital assistant, palmtop computer, cellular radiotelephone, or pager, and the like, and/or the CPE may include one or more of those devices in addition to the devices 1 and 2 shown in Fig. 1. Moreover, the total number and variety of user communication terminals that may be included in the CPE and the overall communication system 10 in general can vary widely, depending on user support requirements, geographic locations, and applicable design/system operating criteria, etc., and are not limited to those depicted in Fig. 1. In general, the teaching of this invention may be employed in conjunction with any suitable types of communication terminals that are capable of communicating with a communication system/network that communicates in accordance with a communication

protocol, such as TCP/IP. It should thus be clear that the teaching of this invention is not to be construed as being limited for use with any particular type of communication terminal or communication protocol.

It also should be noted that, although the equipment 3 and 7 is shown in Fig. 1 as being separate components of the CPE 18 and station 8, respectively, in other embodiments, that equipment may be located within other components of the CPE 18 and station 8. For example, the modem 3 may be located internally within the housing of the user communication terminal 1, and/or the modem 7 may form an integral part of one of the other components of switching station 8.

Referring again to Fig. 1, the components of the central office switching station 8 will now be described. According to a preferred embodiment, the switching station 8 comprises the transceiving equipment 7 and the multiplexer/demultiplexer device 9. Like the equipment 3, the transceiving equipment is preferably an ADSL modem (e.g., ATU-C), and operates in a similar manner as the equipment 3 described above, by demodulating information received from the CPE 18 through link 6 (using, e.g., a known ADSL demodulation technique), and by separating voice

signals from data in the received information, using, for example, an associated splitting device (e.g., a POTS splitter) (not shown). The separated data is then forwarded to the multiplexer/demultiplexer device 9 via a link 10-1, and the separated voice signals are forwarded through a communication link 11 to a PSTN (not shown) for subsequent transmission to a particular receiving destination. The modem 7 also operates by modulating voice signals and data received over the respective links 11 and 10-1, using a known ADSL modulation technique, and by forwarding resulting modulated information to the CPE 18 by way of interface 6.

The multiplexer/demultiplexer device 9 of switching station 8 preferably includes a Digital Subscriber Line Access Multiplixer/Demultiplexer (DSLAM), although in other embodiments, such as those not employing ADSL technology, any other suitable type of multiplexer/demultiplexer device may also be employed. In the preferred embodiment, the multiplexer/demultiplexer 6 operates by coupling signals received over links 10-1 to 10-n onto the communication interface 12, using a known multiplexing technique. The coupled signals are then transmitted to the network 13 over that interface 12. The device 9 also operates by demultiplexing signals received

from the communication interface 12, using a known demultiplexing technique, and by forwarding resulting demultiplexed signals through respective ones of the links 10-1 to 10-n to respective predetermined destinations.

The communication network 13 shown in Fig. 1 will now be described. The communication network 13 preferably includes one or more switches 13a-13n that are interconnected by high-speed optical links, such as, e.g., OC-3 (Optical Carrier) links or other types of high speed links such as, e.g., T-1, T-3, DS-1, or DS-3 links, etc. The switches 13a-13n collectively operate in a known manner by routing information received from the individual communication interfaces 12 and 14 to intended destinations outside of the network 13, based on address information (e.g., IP address information) included in the received information. Preferably, the network 13 operates in accordance with Frame Relay (FR) technology, although in other embodiments, other suitable types of techniques for routing data between particular source and destination points may also be employed, such as Asynchronous Transfer Mode (ATM) technology. As is known in the art, Frame Relay technology is a packet-switching protocol for transmitting intermittent traffic between networks (e.g., LANs or WANs) and

between end points in a communication system. Frame Relay technology places data in a frame for transmission, and provides a permanent virtual circuit (PVC) connection (i.e., a continuous, dedicated connection) between communicating end-points. ATM technology is a network connection switching technology based on transferring data in cells or packets having a fixed size. Individual cells are processed asynchronously relative to other related cells, and are queued before being multiplexed over a transmission path.

Referring to block 17 of Fig. 1, the Internet 17 will now be described. As used herein, the term " Internet" refers to an infrastructure having protocols and operating rules which effectively permit the creation of a world-wide " network of networks" . By connecting a communication device, such as the communication terminal 1 or 22, to the Internet 17, information may be exchanged between that device and any other source/destination device which also is connected to the Internet 17. Thus, a matrix of interconnected communication devices is provided for enabling information to be exchanged between those devices. In general, devices connected to the Internet adhere to TCP/IP protocols.

Traditionally, various types of interconnecting equipment may form the interface 20, for connecting the network 13 to the Internet 17, such as, for example, optical fibers, wires, cables, switches, routers, and other types of communication equipment, although, for convenience, only the links 14 and 16 and the node 15 coupled to the test node 22 are shown in Fig. 1. Preferably, the node 15 (to which test node 22 is coupled) is a router located at a Point of Presence (POP) 15', and no other routers are coupled in the interface 20 between the router 15 and the network 13. A POP generally is provided and maintained by an enterprise, such as an Internet Service Provider (ISP), and is a location at which a network (e.g., network 13) can be connected to (interfaced with) another network entity, such as Internet 17.

Having described the various components of the communication system 10 in detail, an aspect of this invention will now be described. In accordance with this aspect of the invention, the inventor has invented a novel method, apparatus, and program for determining an amount of bandwidth that is available in at least one communication path which couples together nodes in a communication system. The method is preferably performed by exercising the

communication path (formed by the communication system components coupling together the nodes 22 and 1) using information signals, to determine a minimum amount of time it takes for the information signals to traverse the path, in each direction, and by performing a predetermined algorithm employing the determined amount of time to calculate the bandwidth (in one of those directions). Preferably, the exercising operation includes a step of using first information signals to exercise a first portion of the communication path, formed by the links 24a, 24b and the router 15, to determine an amount of queuing delay (QD) in the router 15, based on a first predetermined algorithm. The exercising operation preferably also includes another step of using second information signals to exercise a second, larger portion of the communication path, coupling the test node 22 to the user communication terminal 1, to determine the minimum amount of time (also referred to as a " round trip time $RTT_{T\text{-}CPE}$") it takes for the information signals to be transferred bidirectionally between the test node 22 and user communication terminal 1 by way of that path. The bandwidth in question (e.g., the downlink bandwidth available in the portion of the communication path formed by the components coupling the user communication terminal 1 to the router 15),

is then estimated based a second predetermined algorithm defining the bandwidth in terms of the queuing delay of the router 15 and the determined round trip time RTT_{T-CPE}.

Before describing the method of the invention in detail, the derivation of the first and second predetermined algorithms will first be described. As was described above, the first predetermined algorithm is employed for calculating the approximate amount of queuing delay (QD) of the router 15. The algorithm has a preliminary form defined in terms of the following basic formula (F1):

$$RTT_{IF1-IF2} = QD + PTT \qquad \text{(F1)}$$

wherein $RTT_{IF1-IF2}$ represents the amount of time it takes for a hypothetical information packet originally transmitted by the test node controller 21a (through interface (IF1)) to the router 15, to be returned to that controller 21a by the router 15 (through test node interface (IF2)). The term PTT of formula (F1) represents a packet travel time, and is defined by the following relationship (F2):

$$PTT = T_{IF1-POP} + T_{POP-IF2} \qquad \text{(F2)}$$

In formula (F2), $T_{IF1-POP}$ represents the amount of time it takes for a packet to travel from the test node

controller 21a to the router 15 by way of the interface (IF1) and link 24a, and $T_{POP-IF2}$ represents the amount of time it takes for a packet to travel from the router 15 to the test node controller 21a by way of the link 24b and interface (IF2). Terms $T_{IF1-POP}$ and $T_{POP-IF2}$ of formula (F2) may also be expressed in terms of a relationship defined by the following formulas (F3) and (F4), respectively:

$$T_{IF1-POP} = Bits_{IF1-POP} / BW_{T-POP} \quad (F3)$$

$$T_{POP-IF2} = Bits_{POP-IF2} / BW_{T-POP} \quad (F4)$$

wherein $Bits_{IF1-POP}$ represents the number of bits of a hypothetical information packet transmitted from the test node 22 to the router 15 by way of interface (IF1) and link 24a, $Bits_{POP-IF2}$ represents the number of bits of a hypothetical packet returned by router 15 to the test node 22 by way of link 24b and interface (IF2), and BW_{T-POP} represents the amount of bandwidth capacity provided by each link 24a, 24b coupled between the test node 22 and router 15. Assuming that the number of bits included in the hypothetical information packet transmitted from the test node 22 to the router 15 is the same as the number of bits included in the hypothetical returned packet, then, by substituting the right side of each formula (F3) and (F4) for the terms $T_{IF1-POP}$ and $T_{POP-IF2}$, respectively, in

the above formula (F2), and then simplifying the formula (F2), the following simplified formula (F5) can be obtained:

$$PTT = 2*(PS) / BW_{T\text{-}POP} \qquad (F5)$$

wherein (PS) represents the size (in bits) of each hypothetical information packet, and, as was previously described, $BW_{T\text{-}POP}$ represents the amount of bandwidth capacity provided in each link 24a, 24b coupled between the test node 22 and router 15. Now, by substituting the right side of the formula (F5) for the term PTT appearing in formula (F1) above, and then solving the resulting formula (F1) for the term (QD) (defining the theoretical queuing delay of router 15), the following formula (F6) can be obtained, which represents the first predetermined algorithm referred to above:

$$QD = RTT_{IF1\text{-}IF2} - (2*(PS) / BW_{T\text{-}POP}) \quad (F6)$$

In formula (F6), and as was previously described, $RTT_{IF1\text{-}IF2}$ represents an approximation of the amount of time it takes for a hypothetical information packet originally transmitted from test node controller 21a (through interface (IF1)) to the router 15, to be returned to that controller 21a by router 15 (through

test node interface (IF2)), and $BW_{T\text{-}POP}$ and (PS) represent the same information as described above.

Having described the manner in which the first predetermined algorithm (F6) is derived, the manner in which the second predetermined algorithm is derived will now be described. As was previously described, the second predetermined algorithm is employed for calculating the amount of downlink bandwidth available in a communication path formed by the portion of the communication system 10 coupled between the user communication terminal 1 and router 15. That algorithm may be derived based on the following preliminary relationship (F7):

$$RTT_{T\text{-}CPE} = T_{IF1\text{-}POP} + MQD + T_{POP\text{-}CPE} + TR_{CPE\text{-}POP} + MQD + TR_{POP\text{-}IF2} \quad (F7)$$

wherein $RTT_{T\text{-}CPE}$ represents the amount of time it takes for a second, return hypothetical information packet to be received by the test node controller 21a from user communication terminal 1, relative to a time when a first hypothetical information packet is transmitted from the test node controller 21a to that terminal 1. As was previously described, the term $T_{IF1\text{-}POP}$ in formula (F7) represents the amount of time it takes for the first hypothetical information packet to travel from the test node controller 21a (through interface (IF1)) to the router 15, the term MQD represents an estimated

minimum queuing delay of the router 15 (which is determined as described below), and the term $T_{POP-CPE}$ represents the amount of time it takes for the first hypothetical information packet to travel from the router 15 to the controller 21a of user communication terminal 1. Moreover, the term $TR_{CPE-POP}$ represents a theoretical amount of time it takes for the second, return hypothetical information packet to be received by the router 15, after the first hypothetical information packet is received by the controller 21a of terminal 1 (although the signal delay within the user communication terminal 1 is typically negligible), and the term $TR_{POP-IF2}$ represents the amount of time it takes for the second, return hypothetical information packet to travel from the router 15 to the controller 21a of test node 22.

Based on a known relationship between an information packet size and the bandwidth of a communication path transmitting the packet, the terms $T_{IF1-POP}$ and $T_{POP-CPE}$ in formula (F7) can be substituted for, and the two terms MQD in that formula (F7) can be combined to provide following formula (F8):

$$RTT_{T-CPE} = ((PS)/BW_{T-POP}) + 2*MQD + ((PS)/BW_{POP-CPE}) + TR_{CPE-POP} + TR_{POP-IF2} \quad (F8)$$

wherein (PS) represents the size of the first hypothetical information packet, $BW_{T\text{-}POP}$ represents an amount of bandwidth available in each individual link 24a and 24b coupled between the test node 22 and router 15, $BW_{POP\text{-}CPE}$ represents an amount of downlink bandwidth available in the communication path portion formed by the portion of the communication system 10 interposed between router 15 and user communication terminal 1, and the terms MQD, $TR_{CPE\text{-}POP}$, and $TR_{POP\text{-}IF2}$ represent the same information as was described above.

Assuming that the value of the term (PS) is substantially greater than the value of each term $TR_{CPE\text{-}POP}$ and $TR_{POP\text{-}IF2}$ (i.e., the size of the first hypothetical information packet is substantially greater than that of the second, return hypothetical information packet), then the effect of the values represented by terms $TR_{CPE\text{-}POP}$ and $TR_{POP\text{-}IF2}$ in the overall formula (F8) is negligible, and can be ignored. As a result, the formula (F8) can be simplified to provide the following formula (F9):

$$RTT_{T\text{-}CPE} = ((PS)/BW_{T\text{-}POP})+2*MQD+((PS)/BW_{POP\text{-}CPE}) \qquad \text{(F9)}.$$

By rearranging the terms of that formula (F9) and solving for $BW_{POP\text{-}CPE}$ (representing the available bandwidth in the communication path portion coupling the user communication terminal 1 to router 15), the

following formula (F10) can be obtained, which represents the second predetermined algorithm (F10) referred to above:

$$BW_{POP\text{-}CPE} = (PS)/(RTT_{T\text{-}CPE} - ((PS)/BW_{T\text{-}POP})-2*MQD) \quad (F10).$$

In the second predetermined algorithm (F10), predetermined values representing an information packet size and a bandwidth (available in each individual link 24a, 24b), respectively, may be substituted for the terms (PS) and $BW_{T\text{-}POP}$ (i.e., the values of those terms are known, as will be described below), leaving $RTT_{T\text{-}CPE}$ and MQD as the only unknown variables included in the algorithm (F10). In accordance with this invention, values for those unknown variables are determined using a novel method of this invention, and the second predetermined algorithm is solved to determine the amount of downlink bandwidth ($BW_{POP\text{-}CPE}$) available in the communication path portion coupled between the user communication terminal 1 and the router 15. The manner in which the method of the invention is performed, and the manner in which the first and second predetermined algorithms are employed in the invention, will now be described in detail, with reference being made to the flow diagram depicted in Figs. 3a-3c.

In accordance with a preferred embodiment of the invention, the method is performed in two stages. A first stage includes the steps depicted in Fig. 3a, and is performed to approximate a minimum amount of queuing delay (MQD) of the router 15. A second stage of the method includes the steps shown in Figs. 3b and 3c, and is performed to determine a value for the term $RTT_{T\text{-}CPE}$ described above, and also to determine the amount of downlink bandwidth ($BW_{POP\text{-}CPE}$) available between the router 15 and user communication terminal 1.

In step A1 of Fig. 3a, the first stage of the method is started, and it is assumed that the user communication terminal 1 is "connected" to the Internet 17 through the intermediate components 5, 3, 6, 7, 10-1, 9, 12, 13, and 20 of the communication system 10, and that the test node 22 also is "connected" to the Internet 17 through components (IF1), (IF2), 24a, 24b, 15, and 16 of the system 10. For example, each terminal 1 and 22 may be connected to the Internet 17 in response to a user of the terminal causing a predetermined icon presented on the display 21e of the terminal to be selected, in which case one of the programs stored in the memory 21c of the terminal responds by communicating through the corresponding intermediate components to connect the

terminal to the Internet 17, in accordance with known TCP/IP protocols.

In step A2, it is assumed that the user of test node 22 operates the user interface 21d of that test node 22 to cause a predetermined view (not shown) to be presented on the display 21e. Preferably, the predetermined view prompts the user to specify an address (e.g., an IP address) of one of the interfaces (IF1) and (IF2) to which the user desires information to be sent. Assuming that the user then operates the user interface 21d to enter into the controller 21a information specifying the address of the interface (IF2) (i.e., NIC2) of test node 22, and then enters a command specifying that information packets be transmitted to that destination, then the controller 21a responds by transmitting an information packet to the router 15 by way of interface (IF1) and communication link 24a (step A3). Preferably, the transmitted information packet has a format in accordance with, for example, RFC 791 (or later revisions thereof), and has a size that is predetermined based on a value of a Packet Size Variable (PSV1) 36 stored in the memory 21e of the test node 22, although in other embodiments, that value may be specified in step A3 by the user through user interface 21d (in response to, for example, a

prompt being presented on display 21e in step A2). The transmitted information packet preferably includes the user-specified address (i.e., the address of (IF2)), within, for example, a Destination IP Address field 32 in a header 30 of the packet (see, e.g., Fig. 4), and information identifying the time at which the packet was transmitted from the controller 21a in step A3. That packet transmission time information is preferably included in a data field 34 appended to the header field 30 (Fig. 4), and may be determined using any known technique. For example, the controller 21a may refer to an internal clock 21f within the controller 21a, immediately prior to transmitting the packet, to determine the packet transmission time.

At some time after receiving the information packet transmitted by test node 22 to the router 15 in step A3, the router 15 performs a known routing process to correlate the destination address from field 32 of the received packet to corresponding information stored in an internal routing table (not shown) of the router 15. Based on that corresponding information, the router 15 then forwards the packet through an output port specified by the corresponding information (which, in this case, specifies the output coupled to the link 24b) in the routing table. As a result, the packet is returned to the controller 21a of the test

node 22, by way of the link 24b and the interface (IF2) (step A4).

In step A5, the controller 21a of test node 22 responds to receiving the information packet returned by the router 15 in step A4 by referring to the internal clock 21f to determine the receipt time of the packet (i.e., the time at which the packet is received). Then, in step A6 the controller 21a extracts the information specifying the packet's original transmission time from field 34 of the received packet, and employs that information and the packet receipt time determined in step A5 to determine a difference between the packet receipt time and the original packet transmission time. For example, the controller 21a may determine that difference by subtracting the original packet transmission time from the packet receipt time. The difference value determined in step A6 represents the amount of time taken for the packet transmitted by the test node controller 21a in earlier step A3, to be returned to that controller 21a by the router 15 in step A4.

After step A6 is performed, control passes to step A7 where the controller 21a of test node 22 substitutes the difference value determined in previous step A6, the value specified by the PSV1 variable 36 stored in

memory 21c, and the value of a variable (VBW_{T-POP}) 37 (stored in memory 21c of test node 22) representing the amount of bandwidth available in the link 24a, into the first predetermined algorithm (F6) described above, in place of the terms $RTT_{IF1-IF2}$, (PS), and BW_{t-POP}, respectively, in that algorithm (F6). The controller 21a then solves the algorithm (F6), which is reproduced below for convenience, to determine a value of the term (QD).

$$QD = RTT_{IF1-IF2} - (2*(PS) / BW_{T-POP}) \quad (F6)$$

The value of (QD) determined as a result of the performance of the algorithm (F6) represents an approximation of the amount of propogation delay experienced by the information packet while passing through the router 15, as a result of the queuing delay in the router 15, and is stored in the memory 21c of the test node 22 by the controller 21a of that test node 22 (step A7).

Thereafter, control passes to step A8 where the controller 21a of the test node 22 determines whether or not a predetermined number of information packets have been transmitted by the test node 22 to the router 15, since the method began in earlier step A1. For example, the controller 21a may perform step A8 by comparing a value of a counter variable (not shown)

indicating the number of packets already transmitted by the test node 22, to a predetermined value (not shown), to determine whether or not the value of the counter variable equals the predetermined value. Preferably, the predetermined value is large enough for enabling a large number of router queuing delay samples to be obtained.

If the performance of step A8 results in a determination of 'No' ('N' at step A8), then control passes back to step A2 where the method then continues in the above-described manner. If, on the other hand, the performance of step A8 results in a determination of 'Yes' ('Y' at step A8), then control passes to step A9, where a further step is performed.

According to a preferred embodiment of the invention, in step A9 the controller 21a examines all of the queuing delay (QD) values stored previously in the memory 21c during previous performances of step A7, to determine which one of those values is smallest, and then stores the determined smallest value in the memory 21c of the test node 22 (step A10). That value represents the minimum queuing delay of the router 15 determined during the performance of the first stage of the method of the invention.

It should be noted that any suitable, known technique may be employed by the controller 21a to determine the smallest queuing delay value in step A9 (such as, e.g., a technique comparing pairs of the values to determine a smallest value), and thus that step will not be described in further detail herein. In other embodiments of the invention, step A9 may be performed using known techniques to determine a median value, average value, or other desired value among the (QD) values stored in the memory 21c during previous performances of step A7, depending on applicable operating criteria.

After step A10 is performed, control passes through connector (A) to Fig. 3b, where the second stage of the method is started. In step A11 of Fig. 3b, the controller 21a of test node 22 presents a view (not shown) on the test node display 21e prompting the user to specify an address (e.g., an IP address) of a destination for which he desires to conduct an Internet connection bandwidth test. Assuming that the user desires to conduct such a test for the user communication terminal 1, and thus enters information specifying an address of that terminal 1 into controller 21a in the above-described manner, then the controller 21a responds in step A12 by transmitting an information packet to the router 15 by way of test

node interface (IF1) and communication link 24a (step A12). The information packet transmitted in step A12 preferably has a message format in accordance with, for example, RFC 791 (or later revisions thereof), and includes various data fields, such as, for example, those depicted in Fig. 4. Preferably, data field 34 (Fig. 4) of the transmitted information packet includes information specifying the time at which the packet was transmitted from the controller 21a in step A12 (as determined by internal clock 21f), a source address field 40 of the packet includes information identifying a source address of the interface (IF2) (i.e., NIC2), a Destination IP Address field 32 of the packet includes information identifying the destination address specified by the user in previous step A11, and a Time-To-Live field 39 in a header 30 of the packet includes information specifying a maximum number of hop counts through which the information packet is permitted to travel. That number of hop counts preferably equals the number of hops existing in the portion of the communication system 10 provided from the test node 22 to the user communication terminal 1, and, in this example, is '2' (e.g., the router 15 represents one hop, and the user communication terminal 1 represents another hop).

At some time after receiving the information packet transmitted by the test node 22 in step A12, the router 15 extracts the hop counts value (e.g., '2') from the Time-To-Live field 39 in the received packet, reduces that value by '1', and then reinserts the resulting reduced value (e.g. '1') back into the field 39 of the received packet (step A13). Also in step A13, the router 15 then operates in the above-described manner to correlate the destination address from field 32 of the received packet to corresponding information stored in the internal routing table (not shown) of the router 15. Then, based on that corresponding information, the router 15 forwards the packet, including the reduced hop counts value, through the output specified by the corresponding information (which, in this case, specifies the output coupled to communication link 14) in the routing table. As a result, the information packet is forwarded through the link 14 to the network 13, which, in turn, responds in step A14 by operating in the above-described manner to forward the packet to the user communication terminal 1 (by way of system components 12, 8, 6, 3, and 5), based on the destination address included in field 32 of the packet (step A14).

In step A15 of Fig. 3b, the controller 21a of user communication terminal 1 responds to receiving the information packet in previous step A14 by extracting the hop counts value (e.g., '1') from the Time-To-Live field 39 of the received packet, and by then reducing the extracted value by '1' to obtain a resulting value of '0'. The controller 21a of user communication terminal 1 then recognizes that the obtained value is '0' in step A16, and then responds by extracting (i) the information specifying the packet transmission time from data field 34 of the received information packet, and (ii) the source address information from source address field 40 of the received packet, and by forming an error message that includes the extracted information (i) and (ii) (step A17). Preferably, the message formed in step A17 is an Internet Control Message Protocol (ICMP) message having a message format in accordance with, for example, RFC 792 (or later revisions thereof), includes data fields as shown in, for example, Fig. 5, and has a size (e.g., 36 bytes) that is substantially less than that of the information packet previously transmitted by the test node 22 in step A12 (to render the terms $TR_{CPE-POP}$ and $TR_{POP-IF2}$ in formula (F8) negligible). In the preferred embodiment, the information (i) extracted in step A17 is included in a beginning portion (e.g., the first

eight bytes) of a data field 44 of the error message, and the information (ii) extracted in step A17 is included in a Destination Address field 46 of a header 42 in the message.

After forming the error message in the above-described manner, the controller 21a of user communication terminal 1 transmits the message as an information packet, to the network 13, via system components 5, 3, 6, 8, and 12 (step A18). The network 13 then responds in step A19 in the above-described manner by forwarding the message through link 14 to the router 15, based on the information included in the Destination Address field 46 of the message. Thereafter, the router 15 responds to receiving the message by operating in the above-described manner to cause the received message to be forwarded to the test node controller 21a, by way of components 24b and (IF2), based on the information included in the Destination Address field 46 of the message and corresponding information stored in the internal routing table of router 15 (step A20). Control then passes through connector (B) to step A21 of Fig. 3c.

In step A21 of Fig. 3c, the controller 21a of test node 22 responds to receiving the message routed thereto in previous step A20 by referring to the

internal clock 21f to determine the receipt time of the message in the controller 21a. Then, in step A22, the controller 21a extracts the information specifying the original packet transmission time from field 44 of the received message, and employs that information and the message receipt time determined in step A21, to determine a difference between the message receipt time and the original packet transmission time (step A22), in a similar manner as was described above. The value of the difference determined in step A22 also is referred to in this description as a " round-trip time ($RTT_{T\text{-}CPE}$) value" , and represents an approximation of the amount of time taken for the test node controller 21a to receive the return error message (from user communication terminal 1), after originally transmitting the information packet in earlier step A12 (Fig. 3b). That determined value is then stored by the controller 21a of the test node 22 in the test node memory 21c.

After step A22 is performed, control passes to step A23 where the test node controller 21a determines in the above-described manner whether or not a predetermined number of information packets has been transmitted by the test node 22 to the user communication terminal 1, since the second stage of the method began in earlier step A11. Like step A8 of

the first method stage described above, step A23 is preferably performed so that a large number of round-trip time samples are obtained.

If the performance of step A23 results in a determination of 'No' ('N' at step A23), then control passes through connector (C), back to step A12 of Fig. 3b, where the method then continues in the above-described manner. If, on the other hand, the performance of step A23 results in a determination of 'Yes' ('Y' at step A23), then control passes to step A24 where the test node controller 21a examines the RTT_{T-CPE} values stored previously in the test node memory 21c during the previous performances of step A22, to determine which one of those values is smallest, and then stores the determined smallest value in the memory 21c of the test node 22 (step A24). That smallest value represents the minimum determined amount of time taken (i.e., the minimum round-trip time) for the test node controller 21a to receive an error message from the user communication terminal 1, after transmitting a corresponding, error-provoking information packet in step A12 during the performance of the second stage of the method.

It should be noted that, as for the minimum queuing delay determination described above with respect to

the first method stage (Fig. 3a), any suitable technique may be employed by the test node controller 21a to determine the smallest round-trip time (RTT_{T-CPE}) value in step A24, and one skilled in the art would readily appreciate in view of this description how to formulate an algorithm for use by the controller 21a in making such a determination. In other embodiments of this invention, step A24 may be performed to determine a median value, average value, or other desired value among the determined round-trip time (RTT_{T-CPE}) values stored in memory 21c, depending on applicable performance criteria.

After the minimum round-trip time value is determined in step A24, the controller 21a retrieves the values of the variables PSV2 and VBW_{T-POP} and the minimum queuing delay (MQD) value (stored in earlier step A10) from the test node memory 21c, and substitutes the retrieved values into the second predetermined algorithm (F10) in place of the terms (PS), BW_{T-POP}, and MQD, respectively, in that algorithm. The controller 21a also substitutes the minimum round-trip time value determined in previous step A24 into the second predetermined algorithm (F10), in place of the term RTT_{T-CPE} in that algorithm. Thereafter, the test node controller 21a performs the second predetermined algorithm (F10), which is reproduced below for

convenience, to solve for the term $BW_{POP-CPE}$ in that algorithm (step A25).

$$BW_{POP-CPE} = (PS)/(RTT_{T-CPE} - ((PS)/BW_{T-POP})-2*MQD) \quad (F10).$$

The value obtained as a result of the performance of the algorithm (F10) in step A25 represents an approximation of the maximum amount of downlink bandwidth available in the communication path portion formed by the intermediate system components 5, 3, 6, 7, 10-1, 9, 12, 13, and 14, coupled between the user communication terminal 1 and router 15.

After the downlink bandwidth value is determined in step A25, control passes to step A26, where the controller 21a of test node 22 causes that value to be displayed on the display 21e of the node 22. In other embodiments of the invention, the controller 21a may also cause the value to be stored in the test node memory 21c, and/or forwarded in a message to the user communication terminal 1 or some other predetermined destination (not shown) (step A26), where, the value may be presented to another user or stored for later retrieval thereof. Thereafter, control passes to step A27 where the method terminates.

Another embodiment of this invention will now be described, with reference being made to Fig. 6. In

Fig. 6, a communication system 50 constructed in accordance with this embodiment is shown. The system 50 comprises the same components 1 to 22 as the system 10 of Fig. 1 described above. However, in this embodiment, the test node 22 is coupled to the router 15 through the network 13 and a plurality of bidirectional communication links 24a-1, 24a-2, 24b-1, and 24b-2. Preferably, the links 24a-1 and 24b-1 bidirectionally couple the respective test node interfaces (IF1) and (IF2) to a switch 13n of the network 13, and the links 24a-2 and 24b-2 each bidirectionally couple the switch 13n to the router 15, although in other embodiments, the links 24a-1, 24b-1 and 24a-2, 24b-2 may be coupled to one or more different ones of the switches 13a-13n in the network 13 and the test node 22 may be coupled to the router 15 through more than one of those switches 13a-13n. Preferably, the links 24a-1, 24a-2, 24b-1, and 24b-2 are high speed T1 or T3 links. The links 24a-1, 24a-2, 24b-1, and 24b-2 provide permanent virtual circuits (PVCs) between the test node 22 and router 15, in a case where network 13 is a Frame Relay network, and provide virtual circuits (VCs) between the test node 22 and router 15, in a case where the network 13 is an ATM network.

In accordance with this embodiment of the invention, the above-described method of Figs. 3a-3c is performed in a similar manner as was described above, except that information packets exchanged between the test node 22 and router 15 are transferred through the network 13 (by way of links 24a-1, 24a-1 or 24b-1, 24b-2), and information packets that are forwarded from the test node 22 to the user communication terminal 1 pass through the components 24a-1, 24a-2 (or 24b-1, 24b-2) 15, 14, 13, 12, 8, 6, 3, and 5, in that order. Conversely, information packets that are forwarded from the terminal 1 to the test node 22 pass through the same components, but in a reversed order.

A method for determining the amount of bandwidth available in a communication path in accordance with a further embodiment of this invention will now be described. The method according to this embodiment of the invention is performed by transferring a file between nodes that are coupled together through the path, and by determining the rate at which the file is received at the receiving node, to obtain the communication path bandwidth. The method of this embodiment of the invention may be employed in conjunction with either of the system configurations depicted in Figs. 1 and 6, or in any other suitable type of communication system/network, although for

convenience, the following description is described only in the context of the method being employed in the system 10 of Fig. 1.

Referring now to Fig. 7, a flow diagram of the method according to this embodiment of the invention is shown. In step A100 of Fig. 7 the method is started, and it is assumed that both the user terminal 1 and the test node 22 are connected to the Internet 17 in the above-described manner. It also is assumed that the user of user communication terminal 1 desires to determine the uplink and downlink connection bandwidths provided by the communication system 10 for the terminal 1, and thus operates the user interface 21d of that the terminal 1 to cause a command to be entered into the controller 21a requesting that the user be notified of those bandwidths (step A102).

Thereafter, in step A103 the controller 21a of user communication terminal 1 responds to receiving the command by forming a message that includes information specifying the user request, the address of user communication terminal 1 (representing a source address), and the destination address of a predetermined one of the interfaces (IF1), (IF2) of test node 22, such as, e.g., the interface (IF1). The controller 21a of terminal 1 then communicates the

formed message to the test node 22, by way of intermediate system components 5, 3, 6, 8, 12, 13, 14, and 15, and 24a, in the above-described manner (step A103). In response to eventually receiving that message, the controller 21a of test node 22 responds by extracting the source and destination address information from the message and by retrieving a predetermined file from the memory 21c of the test node 22. Preferably, the predetermined file has a size that is substantially larger than the amount of the downlink bandwidth (e.g., on the order of about 100 times the amount of the downlink bandwidth) expected to be available in the communication path formed by the system components coupled between the node 22 and the terminal 1. That file also preferably includes a first predetermined code (e.g., a Start-of-File code) at a beginning portion of the file, and a second predetermined code (e.g., an End-of-File code) at an end portion of the file (see, e.g., RFC 959).

Thereafter, the test node controller 21a downloads the retrieved file, along with the extracted source and destination address information, to the user communication terminal 1 by way of system components 24a, 15, 14, 13, 12, 8, 6, 3, and 5 (step A104). Preferably, the downloading step A104 is performed in accordance with, for example, RFC 959 (File Transfer

Protocol) (or later revisions thereof), and the downloaded information has a format in accordance with that protocol, although in other embodiments, any other suitable types file transfer protocols/message formats may also be employed.

In step A105, the controller 21a of the user communication terminal 1 measures the period of time taken for the file to be downloaded into that controller 21a, based on the first and second predetermined codes included in the file and time kept by the internal clock 21f of the terminal 1, and also determines the size of the file. For example, the controller 21a may measure the file download time by detecting the receipt of the first predetermined code included in the received file, and by then referring to the internal clock 21f of that controller 21a to determine the receipt time of that first predetermined code. Also, as the end portion of the file is received, the controller 21a detects the second predetermined code (e.g., End Of File code) included in that end portion of the received file, and again refers to the internal clock 21f to determine the receipt time of the second predetermined code. Thereafter, the controller 21a determines the period of time taken for the file to be downloaded thereto by subtracting the determined receipt time of the first

predetermined code from the determined receipt time of the second predetermined code. Also by example, the controller 21a may determine the size of the downloaded file by setting a predetermined counter variable (initially '0') equal to value '1', in response to detecting a first byte (e.g., Start-of-File) of the file, and by then increasing the value of that variable by '1' each time a next byte of the downloaded file is received, to determine the total number of bytes included in the downloaded file. The value of that counter variable remaining after a last, predetermined byte (e.g., End-Of-File) of the file has been received indicates the size (in bytes) of the downloaded file. Preferably, the controller 21a then multiplies that counter variable value by '8' to determine the total number of bits included in the file (representing the file size in bits), although in other embodiments that step need not be performed.

After determining both the period of time taken for the file to be downloaded into the controller 21a of the user communication terminal 1, and the size of the downloaded file (in step A105), the controller 21a of that terminal 1 then performs a predefined algorithm that employs the determined file size and the time period measured in step A105 to determine a value representing an approximation of the rate at which the

file was downloaded through the communication system 10 (step A106). For example, that algorithm may be performed by dividing the determined size of the downloaded file by the measured download time period. As can be appreciated by one skilled in the art, the value determined in step A106 also represents the amount of downlink bandwidth available in the communication path formed by the intermediate system components 24a, 15, 14, 13, 12, 9, 10-1, 7, 6, 3, and 5, coupled between the node 22 and terminal 1.

After step A106 is performed, control passes to step A107 where the controller 21a of user communication terminal 1 presents the determined downlink bandwidth value to the user of terminal 1 through the output user-interface 21e (step A107). The controller 21a also uploads the received file, along with information identifying the addresses of the respective terminals 1 and 22, back to the test node 22, by way of the intermediate system components 5, 3, 6, 8, 12, 13, 14, 15, 24a, and (IF1), in the above-described manner (step A108). Preferably, that uploading step A108 is performed in accordance with, for example, RFC 959 (File Transfer Protocol) (or later revisions thereof), and the uploaded information has a format in accordance with that protocol, although in other

embodiments, any other suitable types file transfer protocols/message formats may also be employed.

In step A109, the controller 21a of the test node 22 measures both the period of time taken for the file to be uploaded into that controller 21a, and the size of the uploaded file, in a similar manner as was described above. For example, the controller 21a preferably measures the file upload time period by detecting the first and second predetermined codes included in the respective beginning and ending portions of the file, referring to the time kept by the internal clock 21f (within test node 22), upon detecting each code, to determine the receipt time of each respective code, and by subtracting the determined receipt time of the first predetermined code from that of the second predetermined code, to determine the file upload time period. Also by example, the controller 21a preferably measures the size of the uploaded file by detecting each byte of the uploaded file, as it is being received, and by increasing the value of a predetermined counter variable (initially '0') by '1' in response to detecting each individual byte of the file, to determine (count) the total number of bytes included in the uploaded file. The value of that counter variable remaining after a last, predetermined byte

(e.g., End-of-File) of the file has been received indicates the size (in bytes) of the uploaded file. Preferably, the controller 21a then multiplies that remaining counter variable value by '8' to determine the total number of bits included in the file, although in other embodiments that step need not be performed.

After determining the file size and upload time period in step A109, the controller 21a of test node 22 then performs a predefined algorithm that employs the determined file size and upload time period to determine a value representing the rate at which the file was uploaded through the communication system 10 (step A110). That determined value also represents the amount of uplink bandwidth available in the communication path formed by the system components 5, 3, 6, 7, 10-1, 9, 12, 13, 14, 15, 5, and 24a coupled between the terminal 1 and the node 22. As for the predefined algorithm performed within the terminal 1 in earlier step A106, the predefined algorithm performed by the test node controller 21a in step A110 may be performed by, for example, dividing the determined size of the uploaded file by the determined file upload time period.

Thereafter, in step A111 the test node controller 21a forwards information representing the uplink bandwidth value determined in step A110 in a message to the user communication terminal 1, by way of the intermediate system components 15, 14, 13, 12, 8, 6, 3, and 5 (step A111). In other embodiments, the test node controller 21a may also forward that message to another predetermined destination (not shown), store the value in the test node memory 21c for later retrieval, and/or present the value to a user of the test node 22 via the output user-interface 21e.

In step A112, the controller 21a of user communication terminal 1 responds to receiving the message transmitted by the test node 22 in previous step A111 by presenting the determined bandwidth value included in the received message to the user of the terminal 1, through the output user-interface 21e of the terminal 1. In other embodiments, the controller 21a may store that value in the memory 21c of the terminal 1 for later retrieval by the user of that terminal 1, depending on applicable performance criteria. Thereafter, the method terminates.

The foregoing embodiments of the invention enable the bandwidth available in a communication path coupled between nodes in a communication system to be

determined, in a manner which overcomes the problems associated with the prior art methods described above. For example, the method of Figs. 3a-3c can be initiated to determine the downlink bandwidth from a single location (i.e., test node 22), and does not require the use of any additional software in the user communication terminal 1. The methods of the invention may be employed regardless of the type of backbone employed in the communication system (e.g., the methods may be used in FR networks, ATM networks, etc.), and may be employed both in systems in which IP address are statically allocated and systems in which IP addresses are DHCP based.

Also by example, because the test node 22 is coupled to the user communication terminal 1 through the router 15 located at the POP 15', it is not necessary to rebuild any virtual circuits before conducting the methods of the invention, since the router 15 automatically facilitates the transfer of information (e.g., a file or information packets) between those devices, by way of communication path existing between those devices. As a result, problems associated with the rebuilding of virtual circuits are avoided.

It should be noted that while this invention is described in the context of the user communication

terminal 1 communicating with the Internet 17 through a communication system having the particular configurations shown in Fig. 1 and Fig. 6, the invention is not necessarily limited for use only in conjunction with those particular system configurations, but may also be employed in any other suitable types of communication systems/networks, depending on applicable system/network architecture. It also should be noted that although the invention is described in the context of the user communication terminal 1 communicating with the Internet 17 through the central office switching station 8, network 13, and communication interface 20, in other embodiments, no switching station 8, network 13, or interface 20 need be employed, and the user communication terminal 1 may communicate with the Internet 17 through other suitable types of interface components, depending on applicable system architecture. Moreover, although the invention is described in the context of the CPE 18 having the user communication terminal 1 for communicating with other components of the system 10 through the modem 3, in other embodiments, the user communication terminal 1 may be included within a WAN, LAN, or a wireless network, and may communicate with the other components of the system 10 through a network server or wireless transceiver (not shown).

It should be further noted that while the method shown in Figs. 3a-3c is described in the context of the terminal 22 prompting the user to enter destination address information separately in steps A2 and A11, and in the context of the user of terminal 22 initiating the performance of that method, it also is within the scope of this invention for the user of the terminal 22 to enter that information in a single step A2, and/or for the user of terminal 1 to initiate the performance of the method by causing a command to be transmitted to the terminal 22, and/or for the method to be initiated automatically, without any user intervention. Similarly, while the method of Fig. 7 is described in the context of the user of terminal 1 initiating the performance of the method, it also is within the context of this invention for that method to be initiated at the test node 22, either automatically or in response to a user-entered command.

Furthermore, it is within the scope of this invention for the user of either terminal 1, 22 to program the value of one or more of the above-described variables PSV1, PSV2, and VBW_{T-POP} into the controller 21a of test node 22, and for the user of either terminal 1 or 22 to pre-specify the number of packets to be transmitted by the test node 22 during the first and second stages

of the method of Figs. 3a-3c. Furthermore, the sizes (defined by variables PSV1 and PSV2) of the information packets employed in those stages may either be the same or different, depending on applicable performance criteria, user preferences, and the like.

Also, although the method of Figs. 3a-3c is described in the context of there being only the single router 15 coupled between the test node 22 and user communication terminal 1, the invention may also be employed in systems having more or less than that number of routers 15. For example, in cases in which more than a single router exists between the test node 22 and terminal 1, the value included in the TTL field 39 during the performance of the method of Figs. 3a-3c is preferably set to account for that number of routers, and the above-described algorithms preferably are adapted to account for that number of routers (i.e., for enabling a minimum queuing delay (MQD) to be determined for each router, in the above-described manner), in a manner as would be readily appreciated by one skilled in the art in view of this description.

While the invention has been particularly shown and described with respect to preferred embodiments thereof, it will be understood by those skilled in the the

art that changes in form and details may be made therein without departing from the scope and spirit of the invention.